

Distributed Charging Control of Electric Vehicles Using Online Learning

Wann-Jiun Ma, Vijay Gupta, Ufuk Topcu

Abstract—We propose an algorithm for distributed charging control of electric vehicles (EVs) using online learning and online convex optimization. Many distributed charging control algorithms in the literature implicitly assume fast two-way communication between a distribution company and EV customers. This assumption is impractical at present and raises privacy and security concerns. Our algorithm does not use this assumption; however, at the expense of slower convergence to the optimal solution. The proposed algorithm requires one-way communication, which is implemented through the distribution company publishing the pricing profiles of the previous days. We provide convergence results of the algorithm and illustrate the results through numerical examples.

Index Terms—Charging control, demand response, electric vehicles, online convex optimization, online learning, regret minimization

I. INTRODUCTION

Demand response (DR) is an important functionality for the next generation power systems. It empowers the distribution company and its customers to decide collectively, but in a distributed manner, the best way to schedule energy usage. The reader is referred to [7], [11], [22], and references therein for the details of DR. This paper focuses on the flexible load capability offered by electric vehicles (EVs) owned by residential customers.

Large-scale integration of EVs imposes a significant burden on the grid. Particularly, creation of new peaks, peak load amplification [15] and voltage deviations [18] among other effects have been identified as major concerns. To cope with these issues, many algorithms have been proposed to schedule the charging of EVs, e.g., [5], [8], [20], [21], [26]. We are particularly interested in the algorithms such as those proposed in [9], [10], [19], [20], which lead to the convergence of the total load profile to a desired one (for instance, a valley-filling profile) through appropriate price signals transmitted to the owners.

Many of the existing algorithms have analytical convergence guarantees and do not require the customers to share their charging constraints with the distribution company. On the other hand, they require a series of messages to be exchanged among the distribution company and the customers regarding possible price profiles and desired charging profiles in response. As the available power supply and the customer

requirements for charging their EVs change from day to day, these messages need to be exchanged daily to calculate the charging profiles. Since these message exchanges need to be completed before the EVs can begin charging, the algorithms, thus, implicitly assume the presence of a communication infrastructure and protocols that can support low-latency two-way communication between the distribution company and the EV customers. Such infrastructure and protocols have not yet been deployed extensively [2]. Furthermore, the transmitted data carry information about the constraints faced by the individual customer and raises privacy and security concerns.

Motivated by this issue, we propose an online learning and online convex optimization based distributed charging control algorithm. This algorithm requires only one-way communication from the distribution company to the customers. Furthermore, the communication carries information about the pricing profiles of *previous* days.

In our formulation, we model the distribution company and every EV customer as decision makers who wish to optimize their own utility functions. For the distribution company, the payoff is maximized if the total load profile over a day is valley-filling [9], [10]. For the EV customer, the utility function is maximized if the cost to charge the EV over a day is minimized. By designing a suitable pricing policy, the distribution company aims at ensuring that the charging profiles followed by the customers aggregate to a valley-filling profile. Our distributed charging control algorithm is based on an online learning and online convex optimization framework. The only communication that occurs is when the distribution company notifies every customer of the pricing profiles incurred over the previous days. The online learning framework has tremendous popularity in the online convex optimization and machine learning community (see e.g., [4], [6], [13], [24], [27], and references therein). We use a regret minimization algorithm [23] in the online learning framework. The regret minimization algorithm uses a *regret* as the performance measure and provides an iterative way for every decision maker to update its policy such that, at convergence, the policy is optimal in a suitably defined sense. Informally, the regret minimization algorithm operates as follows. Consider a situation in which multiple decision makers need to design their own individual utility functions while satisfying coupled constraints. Furthermore, the utility functions of the decision makers may also be coupled through the actions of multiple decision makers and also the environment. The environment consists of the factors that the decision makers cannot control. The environment is uncertain and time-varying. Decisions are made repeatedly and the resulting payoffs are used to improve the decision policies used by the decision makers. Thus, in every iteration, every decision maker makes a decision given

W.-J. Ma and V. Gupta are with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556, USA. U. Topcu is with the Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104, USA. {wma1, vgupta2}@nd.edu, utopcu@seas.upenn.edu. The work of W.-J. Ma and V. Gupta is supported in part by NSF grant #1239224. The work of U. Topcu is supported in part by NSF grant #1312390. Initial results were presented at IEEE Conference on Decision and Control (CDC), 2014.

the current policy and the realization of the environment. Given the decisions of the decision makers and the realization of the environment, the various decision makers obtain certain payoffs.

Our contributions are two-fold. First, we present a regret minimization based distributed algorithm for charging control of EVs that requires only one-way communication. Second, we allow heterogeneous utility functions for the various EV owners, as would be the case when customers vary in the elasticity of shifting their loads in response to prices.

Some relevant references that apply regret minimization to DR are [14], [16], [25]. In [16], real-time electricity pricing strategies for DR are designed using regret minimization. However, the focus of that work is on optimizing the utility function for the distribution company, and the customer behavior is assumed to be such that the load change is linear in the price variation. The objective of [14] is to design pricing policies for the customers having price responsive loads. The exact demand function of the customer is assumed to be unknown to the pricing policy maker. In [14], the distribution company is the only decision maker, whereas in our work, the distribution company and the EV customers are all decision makers. In [25], regret minimization is used to learn the charging behavior of the EV customers. The price responsiveness for a community of customers is captured through a conditional random field model. The regret minimization algorithm is adopted to learn the parameters of the model.

The remainder of the paper is organized as follows. Section II presents the problem formulation. The basic framework and main results are presented in Section III. Section IV extends the basic framework to some practical charging scenarios. Some numerical examples can be found in Section V. Section VI concludes the paper.

II. PROBLEM FORMULATION

We consider a scenario in which N customers schedule the charging of their electric vehicles (EVs) daily. The charging needs to be completed over a day. Let there be T time slots in a day and denote the set of these time slots by $\mathcal{T} := \{1, \dots, T\}$. Denote the set of EVs by $\mathcal{N} := \{1, \dots, N\}$. Denote the base load on day k by $D^k(t) \in \mathbb{R}$, $t \in \mathcal{T}$. We assume that this base load is unknown to the EV customers and to the distribution company at the beginning of the k -th day when the charging schedules are fixed. Furthermore, the base load may vary from day to day. Denote by $x_i^k(t) \in \mathbb{R}$ the charging rate of the i -th EV in the t -th time slot on the k -th day. The charging profile of the i -th EV on the k -th day is denoted by a vector $x_i^k := (x_i^k(1), x_i^k(2), \dots, x_i^k(T))$. The aggregated charging profile of the EV customers is described by a vector $x^k := (x_1^k, \dots, x_N^k)$. Let $x_i^{\text{low}}(t)$ and $x_i^{\text{up}}(t)$ denote the minimum and maximum charging rates for the i -th EV in the t -th time slot and S_i denote the desired total charge for the i -th EV at the end of the k -th day; thus $S_i = \sum_{t \in \mathcal{T}} x_i^k(t)$, for every k . The total load as seen by the distribution company is the sum of the base load and the charging rates adopted by the EVs.

The objective of the distribution company is to achieve a total load profile that is valley-filling while ensuring that both

the base load and the EVs are supplied with the required amount of energy. The base load is inflexible, while the EV charging profile may be shaped as long as the desired amount of power is provided by the end of the day. The goal of the distribution company can be described as obtaining the aggregated charging profile x^k , for every $k \in \mathbb{N}_{>0}$, that solves

$$\begin{aligned} & \underset{x^k}{\text{minimize}} \quad c_u^k(x^k) \\ & \text{subject to} \quad x_i^{\text{low}}(t) \leq x_i^k(t) \leq x_i^{\text{up}}(t), \quad t \in \mathcal{T}, i \in \mathcal{N}, \\ & \quad \sum_{t \in \mathcal{T}} x_i^k(t) = S_i, \quad i \in \mathcal{N}, \end{aligned} \quad (1)$$

where the distribution company cost function c_u^k is chosen as

$$c_u^k(x^k) := \sum_{t \in \mathcal{T}} (D^k(t) + \sum_{i \in \mathcal{N}} x_i^k(t))^2. \quad (2)$$

The cost function (2) for the distribution company is also considered in [9], [10]. By solving (1), the distribution company can obtain a valley-filling total load profile while satisfying the charging requests of the customers.

To incentivize the customers to choose charging profiles that in aggregate minimize the cost (2), the distribution company designs suitable pricing profiles for the power being supplied to the EVs. Every EV customer fixes the charging schedule at the beginning of the day based on the information about her own constraints and any information provided by the distribution company. A *price-sensitive* EV customer seeks to minimize the total cost of charging by suitably shaping her charging schedule. Thus, the optimization problem for each such customer i , $i \in \mathcal{N}$ is to design a charging profile x_i^k , $k \in \mathbb{N}_{>0}$ that solves

$$\begin{aligned} & \underset{x_i^k}{\text{minimize}} \quad c_i^k(x_i^k, x_{j \neq i}^k) \\ & \text{subject to} \quad x_i^{\text{low}}(t) \leq x_i^k(t) \leq x_i^{\text{up}}(t), \quad t \in \mathcal{T}, \\ & \quad \sum_{t \in \mathcal{T}} x_i^k(t) = S_i, \end{aligned} \quad (3)$$

where c_i^k is a convex function in x_i^k and is also a function of the other customers's charging profiles, where $x_{j \neq i}^k := x_j^k$ for $j \in \mathcal{N}$, $j \neq i$. Since the pricing policy is possibly a function of the base load D^k and other customer's charging profiles, c_i^k inherits these features as well.

The information flow is as follows. The distribution company monitors the total load and publishes the price profile for the previous day as realized according to a fixed and known pricing policy. The customers decide on the charging schedules for the next day with access to these pricing profiles for the previous days. No other communication occurs between the distribution company and the customers, or among the customers.

III. ONLINE LEARNING FRAMEWORK

We now adopt a regret minimization framework to solve both problem (1) and problem (3). The regret minimization framework is used for online learning and optimization and only requires one-way message exchanges between the distribution company and the customers.

Let L be a λ -strongly convex function with respect to a given norm $\|\cdot\|$. Let $D_L(\cdot, \cdot)$ denote the Bregman divergence [23] with respect to L . Let $\|\cdot\|_*$ denote the norm that is dual to $\|\cdot\|$. Let ∇L denote the gradient of L and ∇L^{-1}

denote the inverse mapping of ∇L . For example, if L is the squared Euclidean norm, i.e., $L(\cdot) = \|\cdot\|^2$, then the corresponding Bregman divergence $D_L(x, y)$ is equal to the squared Euclidean distance $\|x - y\|^2$, the inverse mapping $\nabla L^{-1}(x)$ is equal to $\frac{1}{2}x$, and the dual norm is the Euclidean norm.

Customer perspective: For the i -th EV customer, the decision variable is her charging profile x_i^k on the k -th day. The set of feasible charging profiles is given as

$$\mathcal{F}_i := \left\{ x_i^k \in \mathbb{R}^T \mid x_i^{\text{low}}(t) \leq x_i^k(t) \leq x_i^{\text{up}}(t), \right. \\ \left. t \in \mathcal{T}, \sum_{t \in \mathcal{T}} x_i^k(t) = S_i \right\}. \quad (4)$$

We define x_i^* as

$$x_i^* := \arg \min_{x_i \in \mathcal{F}_i} \sum_{k=1}^K c_i^k(x_i). \quad (5)$$

In the regret minimization framework. The notion of regret is used to measure the performance of an online algorithm [24], [27]. For customer $i \in \mathcal{N}$, the customer regret after K days R_i is defined as the difference between the cumulative cost function value of the charging profiles x_i^k , $k = 1, \dots, K$ generated by an online algorithm and the one generated by x_i^* , i.e.,

$$R_i(K, x_i^k) := \sum_{k=1}^K c_i^k(x_i^k) - \min_{x_i \in \mathcal{F}_i} \sum_{k=1}^K c_i^k(x_i). \quad (6)$$

As we mentioned in Section II, c_i^k is possibly a function of the base load D^k and other customer's charging profiles $x_{j \neq i}$, $j \in \mathcal{N}$, $j \neq i$. Since the base load and the charging profiles may change from day to day, c_i^k may not remain the same from day to day a well. While x_i^* remains the same from one day to the next, x_i^* is a suboptimal solution of (3). The regret R_i measures the difference between the performance of the charging profile generated by an online algorithm and the performance that is obtained by the suboptimal charging profile x_i^* . Notice that the suboptimal charging profile x_i^* can only be calculated in hindsight after K days have elapsed.

We adopt the optimistic mirror descent (OMD) algorithm [23] to generate the charging profile update which minimizes the regret (6). On each day, the regret minimization algorithm generates the charging profile update without knowing the current objective function (and its gradient). Specifically, the OMD algorithm iteratively applies the updates

$$h_i^{k+1} = \nabla L_i^{-1}(\nabla L_i(h_i^k) - \eta_i \nabla c_i^k(x_i^k)) \\ x_i^{k+1} = \arg \min_{x_i \in \mathcal{F}_i} \eta_i x_i^T M_i^{k+1} + D_{L_i}(x_i, h_i^{k+1}), \quad (7)$$

where $\eta_i \in \mathbb{R}$ is an algorithm parameter, h_i^k is an intermediate update of the charging profile. For easy of presentation, for the vector $h_i^k \in \mathbb{R}^T$, $L_i(h_i^k)$ is set to $L_i(h_i^k) = \frac{\|h_i^k\|^2}{2}$. The i -th customer may have a prediction M_i^k for the gradient of the cost function c_i^k . For example, on day k , for customer $i \in \mathcal{N}$, one possible option of the prediction M_i^k is the average of the

gradients of the cost functions for the previous days, namely, $M_i^k(x_i^k) = \frac{1}{k-1} \sum_{k=1, \dots, k-1} \nabla c_i^k(x_i^k)$. Intuitively, the iteration (7) updates the charging profile toward the negative gradient direction and projects it onto the set of feasible charging profiles.

As we mentioned in Section II, the customer and the distribution company have different objectives and hence different regrets. We now switch to the distribution company perspective and define the regret minimization framework for the company.

Company perspective: For the distribution company, the decision variable is the aggregated charging profile x^k on the k -th day. The set of the aggregated feasible charging profiles is denoted by $\mathcal{F} := \mathcal{F}_1 \times \mathcal{F}_2 \times \dots \times \mathcal{F}_N$. The distribution company's regret after K days is given by

$$R_u(K, x^k) := \sum_{k=1}^K c_u^k(x^k) - \min_{x \in \mathcal{F}} \sum_{k=1}^K c_u^k(x). \quad (8)$$

We define x^* as

$$x^* := \arg \min_{x \in \mathcal{F}} \sum_{k=1}^K c_u^k(x). \quad (9)$$

The OMD algorithm generates the charging profile update which minimizes the regret (8) as

$$h_u^{k+1} = \nabla L_u^{-1}(\nabla L_u(h_u^k) - \eta_u \nabla c_u^k(x^k)), \\ x^{k+1} = \arg \min_{x \in \mathcal{F}} \eta_u x^T M_u^{k+1} + D_{L_u}(x, h_u^{k+1}), \quad (10)$$

where $\eta_u \in \mathbb{R}$ is an algorithm parameter, h_u^k is an intermediate update of the aggregated charging profile, and M_u^k is the prediction of the gradient of the cost function ∇c_u^k . For easy of presentation, for the vector $h_u^k \in \mathbb{R}^{NT}$, $L_u(h_u^k)$ is set to $L_u(h_u^k) = \frac{\|h_u^k\|^2}{2}$. As an example, the prediction M_u^k can be chosen to be the average of the gradients of the cost functions for the previous days, namely, $M_u^k(x^k) = \frac{1}{k-1} \sum_{k=1, \dots, k-1} \nabla c_u^k(x^k)$.

A. Convergence Results

The following results summarize the convergence of the charging profile updates generated by the OMD algorithm. All proofs can be found in the Appendix.

Proposition III.1. (Convergence of regret): For every $x_i^* \in \mathcal{F}_i$, the iteration (7) converges in the sense that

$$R_i(K, x_i^k) \leq \frac{1}{\eta_i} P_i + \frac{\eta_i}{2} \sum_{k=1}^K \|\nabla c_i^k(x_i^k) - M_i^k\|_*^2, \quad (11)$$

where

$$R_i(K, x_i^k) := \sum_{k=1}^K c_i^k(x_i^k) - \sum_{k=1}^K c_i^k(x_i^*), \\ P_i := \max_{x_i \in \mathcal{F}_i} L_i(x_i) - \min_{x_i \in \mathcal{F}_i} L_i(x_i). \quad (12)$$

In particular, if η_i is chosen as $O(1/\sqrt{K})$, then the average regret, i.e., $R_i(K)/K$, converges to zero as $K \rightarrow \infty$.

Similarly, for every $x^* \in \mathcal{F}$, the iteration (10) converges in the sense that

$$R_u(K, x^k) \leq \frac{1}{\eta_u} P_u + \frac{\eta_u}{2} \sum_{k=1}^K \|\nabla c_u^k(x^k) - M_u^k\|_*^2, \quad (13)$$

where

$$\begin{aligned} R_u(K, x^k) &:= \sum_{k=1}^K c_u^k(x^k) - \sum_{k=1}^K c_u^k(x^*), \\ P_u &:= \max_{x \in \mathcal{F}} L_u(x) - \min_{x \in \mathcal{F}} L_u(x). \end{aligned} \quad (14)$$

In particular, if η_u is chosen as $O(1/\sqrt{K})$, then the average regret, i.e., $R_u(K)/K$, converges to zero as $K \rightarrow \infty$.

Any online algorithm yields a sublinear regret bound as in (11) and (13) is called a no regret algorithm [6], [24]. The results (11) and (13) guarantee that, as the number of days increases, the average performance of the charging profiles generated by the OMD algorithm approaches the performance that is obtained by the charging profiles x^* and x_i^* , $i \in \mathcal{N}$, respectively. Note that while the charging profiles x^* and x_i^* , $i \in \mathcal{N}$ may not solve the optimization problems (1) and (3), respectively, these solutions are optimal for the related problems (5) and (9), respectively. In Section IV-A, we compare the performance of the charging profile generated by the OMD algorithm with the performance that is obtained by the solutions of problems (1) and (3). However, in that case, the convergence guarantees that are obtained are weaker.

B. Design of the Pricing Function

There are no guarantees that the solutions x_i^* , $i \in \mathcal{N}$ of the problem (5) can solve the problem (9). In fact, unless the pricing function c_i^k is carefully designed, these solutions will not be the same since the objectives of the distribution company and the EV customers are different. After some algebraic manipulation of the updates (7) and (10), we observe that the natural choice of c_i^k as

$$c_i^k(x_i^k) = \left(\sum_{j=1}^N x_j^k + D^k \right)^T x_i^k \quad (15)$$

does not lead to the charging profiles (x_1^*, \dots, x_N^*) that reduce the regret of the distribution company to zero.

We now propose a choice of c_i^k to ensure that when each customer minimizes her regret, the aggregated charging profile minimizes the distribution company's regret.

Proposition III.2. *If c_i^k is chosen as*

$$c_i^k(x_i^k) = \left(\frac{1}{2} x_i^k + \sum_{j \neq i}^N x_j^k + D^k \right)^T x_i^k, \quad i \in \mathcal{N}, \quad (16)$$

the customers adopt the iteration (7), and $\eta_u = \frac{1}{2} \eta_i$, then the average regret of the distribution company as defined in (8) converges to zero as the total number of days goes to infinity.

To update the charging profile on day k , the i -th customer needs to know $2x_i^{k-1} + \sum_{j \neq i} x_j^{k-1} + D^{k-1}$ or $\sum_j x_j^{k-1} + D^{k-1}$ depending on whether the pricing function (15) or (16)

is adopted. The distribution company can simply publish the total load information for the previous day. The customers do not need to have full knowledge about how their consumption will map to a corresponding expenditure.

IV. EXTENSIONS

The basic framework presented above can be extended in various directions. They include considering a different definition of regret and incorporating customers who vary in the elasticity of shifting their EV charging load in response to price.

To ensure that the distribution company's average regret has a convergent behavior and the regret is of the order $O(\sqrt{K})$, in the following discussion we assume that the distribution company selects the cost function (16) for each EV customer and sets $\eta_u = \frac{1}{2} \eta_i$, $i \in \mathcal{N}$.

A. Regret with Respect to the Optimal Charging Profiles

The regrets defined in (6) and (8) measure the difference between the performance of the charging profiles generated by our algorithm and the performance that is obtained by the charging profiles x^* and x_i^* , $i \in \mathcal{N}$ that are the solutions of the related optimization problems (5) and (9), respectively. We can instead consider the original optimization problems (1) and (3) to define *tracking regret* after K days as

$$R_u^{\text{tracking}}(K, x^k) := \sum_{k=1}^K c_u^k(x^k) - \min_{x^k \in \mathcal{F}, \forall k \in \mathcal{K}} \sum_{k=1}^K c_u^k(x^k), \quad (17)$$

where $\mathcal{K} := \{1, \dots, K\}$. We define the set $\{x^{k*}, k \in \mathcal{K}\}$ as

$$\left\{ x^{k*} \in \mathbb{R}^{NT}, k \in \mathcal{K} \mid x^{k*} = \arg \min_{x^k \in \mathcal{F}, \forall k \in \mathcal{K}} \sum_{k=1}^K c_u^k(x^k) \right\}. \quad (18)$$

This notion of tracking regret characterizes the difference between the cumulative cost of the charging profiles generated by our algorithm and the cumulative cost of executing the optimal charging profiles that can be calculated only in hindsight. For comparison, we refer to the the regrets (6) and (8) as *static regrets*.

Theorem IV.1. *For every $x^{k*} \in \mathcal{F}$, the OMD algorithm yields that,*

$$\begin{aligned} R_u^{\text{tracking}}(K, x^k) &\leq \frac{1}{\eta_u} \left[L_u(h_u^{K+1}) - L_u(h_u^1) \right] \\ &\quad + \frac{1}{\eta_u} \left[\nabla L_u(h_u^{K+1})^T (x^{K+1*} - h_u^{K+1}) \right. \\ &\quad \left. - \nabla L_u(h_u^1)^T (x^{1*} - h_u^1) \right] \\ &\quad + \frac{1}{\eta_u} \max_{k \in \mathcal{K}} \|\nabla L_u(h_u^k)\| \sum_{k=1}^K \|x^{k*} - x^{k+1*}\| \\ &\quad + \frac{\eta_u}{2} \sum_{k=1}^K \|\nabla c_u^k(x^k) - M_u^k\|_*^2. \end{aligned} \quad (19)$$

In particular, if $\eta_u = O(1/\sqrt{K})$, then the tracking regret is of the order $O(\sqrt{K}[1 + \sum_{k=1}^K \|x^{k*} - x^{k+1*}\|])$.

A comparison of the regret bounds (13) and (19) is of interest. If $\eta_u = O(\sqrt{K})$, in (13), the first term $\frac{1}{\eta_u}P_u$ and the second term $\frac{\eta_u}{2} \sum_{k=1}^K \|\nabla c_u^k(x^k) - M_u^k\|_*^2$ are of order $O(\sqrt{K})$. In (19), the first three terms measure the difference between the initial iterate h_u^1 and the final iterate h_u^{K+1} , the difference between the final iterate h_u^{K+1} and the optimal solution x^{K+1*} , and the difference between the iterate h_u^1 and the optimal solution x^{1*} . The first three terms are of order $O(\sqrt{K})$. The last term in (19), $\frac{\eta_u}{2} \sum_{k=1}^K \|\nabla c_u^k(x^k) - M_u^k\|_*^2$, also appears in (13) and is of order $O(\sqrt{K})$. The fourth term increases as K increases and is of order $O(\sqrt{K} \sum_{k=1}^K \|x^{k*} - x^{k+1*}\|)$. Because of the presence of the fourth term, the regret bound in (19) increases as the variation of the optimal sequence of decisions $\sum_{k=1}^K \|x^{k*} - x^{k+1*}\|$ increases. If the optimal solution remains the same from one day to the next, then $\sum_{k=1}^K \|x^{k*} - x^{k+1*}\| = 0$ and the tracking regret R_u^{tracking} is of the order $O(\sqrt{K})$. On the other hand, if the optimal solution varies significantly from one day to the next, then the tracking regret R_u^{tracking} will be of the order $O(\sqrt{K}[1 + \sum_{k=1}^K \|x^{k*} - x^{k+1*}\|])$, and the average tracking regret will not necessarily converge to zero.

If the distribution company has perfect prediction of the gradient of the cost function, i.e., $\nabla c_u^k(x^k) = M_u^k$ for all k , then the last term in (19) vanishes and the distribution company can set $\eta \rightarrow \infty$ to ensure that the regret bound is zero, i.e., $R_u^{\text{tracking}} \leq 0$. It indicates that the cumulative cost function value of the charging profiles x^k , $k = 1, \dots, K$ generated by the online algorithm (10) and the one generated by the elements in the set $\{x^{k*} \in \mathbb{R}^{NT}, k \in \mathcal{K}\}$ are identical. Note that the elements in the set $\{x^{k*} \in \mathbb{R}^{NT}, k \in \mathcal{K}\}$ solve problem (1).

B. Presence of Inelastic Customers

The discussion so far assumed that all customers were rational in the sense that they wanted to choose their charging profile to solve (3). Furthermore, they were elastic in scheduling their charging (within the constraints pre-specified by $x_i^{\text{low}}(t)$, $x_i^{\text{up}}(t)$, and S_i). We now assume that some customers are either irrational or inelastic and they do not optimize their schedules to solve (3). Suppose that N_l out of N customers are inelastic. Denote the set of inelastic customers by \mathcal{N}_l . For every inelastic customer $i \in \mathcal{N}_l$, we assume that her charging profile remains the same from day to day and is not updated to minimize (3). Equivalently, for the inelastic customers, the cost function c_i^k can be selected as $c_i^k \equiv r$, for all k , where r is an arbitrary constant.

Since the inelastic customers do not carry out any predictions, we set all customers' predictions to zeros, i.e., $M_i^k = 0$, $i \in \mathcal{N}$, $k \in \mathbb{N}_{>0}$. The update of the charging profile for inelastic customer $i \in \mathcal{N}_l$ can thus be written as

$$\begin{aligned} h_i^{k+1} &= \nabla L_i^{-1} \left(\nabla L_i(h_i^k) - \eta_i (\nabla c_i^k(x^k) + \epsilon_i^k) \right), \\ x_i^{k+1} &= \underset{x_i \in \mathcal{F}_i}{\text{argmin}} D_{L_i}(x_i, h_i^{k+1}), \end{aligned} \quad (20)$$

where ϵ_i^k is an error term. For instance, for the cost function (16), the error ϵ_i^k is equal to the total load $\sum_j x_j^k + D^k$ on the k -th day. This error term quantifies the inconsistency between the updates as desired by the distribution company for each customer to execute and the inelastic customer's behavior. Denote by ϵ^k the aggregate error, i.e., $\epsilon^k := (\epsilon_1^k, \dots, \epsilon_N^k)$.

Due to the presence of the inelastic customers, the ability of the aggregated solution to be valley-filling and hence to minimize the cost function in problem (1) is decreased. The performance loss is given in the following result.

Theorem IV.2. Consider that there are N EV customers out of which N_l are inelastic customers. If $\eta_u = \frac{1}{2}\eta_i$, $i \in \mathcal{N}$, then for every $x^* \in \mathcal{F}$, the regret of the distribution company is bounded as

$$\begin{aligned} R_u(K, x^k) &\leq \frac{1}{\eta_u}P_u + \frac{\eta_u}{2} \sum_{k=1}^K \|(\nabla c_u^k(x^k) + \epsilon^k)\|_*^2 \\ &\quad + K \sum_{i \in \mathcal{N}_l} \|F_i\| \|\epsilon_i\|, \end{aligned} \quad (21)$$

where

$$\begin{aligned} P_u &:= \max_{x \in \mathcal{F}} L_u(x) - \min_{x \in \mathcal{F}} L_u(x), \\ \|F_i\| &:= \max_{x, y \in \mathcal{F}_i} \|x - y\|, \quad \|\epsilon_i\| := \max_k \|\epsilon_i^k\|, \quad i \in \mathcal{N}_l. \end{aligned} \quad (22)$$

In (21), the average regret converges to a constant, i.e., $\lim_{K \rightarrow \infty} R_u(K)/K = \sum_{i \in \mathcal{N}_l} \|F_i\| \|\epsilon_i\|$. The size of this constant depends on the error terms ϵ_i , $i \in \mathcal{N}_l$ and the charging constraints of the inelastic customers. The result explicitly quantifies the deviation from the desired performance in terms of the asymptotic average regret of the distribution company. To obtain further insight into the effect of the inelastic customers, we proceed as follows.

Note that the bound (21) depends on the size of the term $\|\epsilon_i\|$. For inelastic customer $i \in \mathcal{N}_l$, the error term $\|\epsilon_i\|$ can be bounded as

$$\begin{aligned} \|\epsilon_i\| &= \max_k \left\| \left(\sum_{i \in \mathcal{N}} x_i^k + D^k \right) \right\| \leq \max_k \left(\sum_{i \in \mathcal{N}} \|x_i^k\| + \|D^k\| \right) \\ &\leq \sum_{i \in \mathcal{N}} \|x_i^{\text{up}}\| + \max_k \|D^k\|. \end{aligned} \quad (23)$$

Assume that the Euclidean ball B is the smallest Euclidean ball containing the set of feasible charging schedules \mathcal{F}_i and r is the radius of the ball B . Then, the term $\|F_i\|$ can be further bounded by $2r$. Since the set \mathcal{F}_i is a polytope, the algorithm in [17] can be adopted to compute this bound efficiently.

Notice that if $N_l = 0$, (21) boils down to (13) (without the prediction, i.e., $M_u^k = 0$ for all $k = 1, \dots, K$). However, it becomes difficult to ensure the feasibility of the problem (1). To improve the ability of the aggregated solution to be valley-filling in the presence of inelastic customers, the distribution company can consider using some loads that can be controlled completely. We now discuss this option.

C. Controllable Customers

On the other end of the spectrum from customers that are completely inelastic are customers that are under the complete

control of the distribution company. The customers under the complete control adopt the charging profiles assigned by the company. The charging constraints of these customers are also known and controlled by the distribution company. In practice, the distribution company can offer a contract to a subset of customers offering a special price to be such a controllable load. This contract-based direct load control has been implemented by many distribution companies, e.g., [1], [3].

To introduce such directly controlled customers in our formulation, we allow the charging constraints of the controllable customers to be relaxed (thus enlarging the set of feasible charging profiles for the controllable customers). Denote the set of controllable customers by \mathcal{N}_c . On day k , for the controllable customer $i \in \mathcal{N}_c$, the set of the relaxed feasible charging profiles $\tilde{\mathcal{F}}_i$ is defined as

$$\tilde{\mathcal{F}}_i := \left\{ \tilde{x}_i^k \in \mathbb{R}^T \mid \tilde{x}_i^{\text{low}}(t) \leq \tilde{x}_i^k(t) \leq \tilde{x}_i^{\text{up}}(t), \right. \\ \left. t \in \mathcal{T}, \sum_{t \in \mathcal{T}} a_i \tilde{x}_i^k(t) = \tilde{S}_i \right\}, \quad (24)$$

where $a_i = \{0, 1\}$, \tilde{x}_i^k is the feasible charging profile in the set $\tilde{\mathcal{F}}_i$, and \tilde{x}_i^{low} , \tilde{x}_i^{up} , \tilde{S}_i , the relaxed minimum charging rate, the relaxed maximum charging rate, and the relaxed total charge, respectively. The aggregated relaxed charging profile of the EV customers is described by a vector \tilde{x}^k which consists of all elements in the set

$$\left\{ \tilde{x}_i^k \in \mathbb{R}^T \mid i \in \mathcal{N}_c \right\} \cup \left\{ x_i^k \in \mathbb{R}^T \mid i \in \mathcal{N}_l, i \in \mathcal{N}_p \right\}, \quad (25)$$

where \mathcal{N}_p denotes the set of the price-sensitive customers.

For controllable customer $i \in \mathcal{N}_c$, we define \tilde{x}_i^* as

$$\tilde{x}_i^* := \arg \min_{\tilde{x}_i \in \tilde{\mathcal{F}}_i} \sum_{k=1}^K c_i^k(\tilde{x}_i). \quad (26)$$

We use $\tilde{\mathcal{F}}$ to denote the set of the feasible aggregated charging profiles including the controllable customers. We define \tilde{x}^* as

$$\tilde{x}^* := \arg \min_{\tilde{x} \in \tilde{\mathcal{F}}} \sum_{k=1}^K c_u^k(\tilde{x}). \quad (27)$$

There are different ways to relax the set of feasible charging profiles for the controllable customers. For example, if we select $a_i = 0$ and $\tilde{S}_i = 0$, then the equality constraint is removed from the set of feasible charging profiles for the i -th customer, $i \in \mathcal{N}_c$, namely, the i -th controllable customer removes her total charging sum requirement daily. We can also relax the charging deadline by adjusting $\tilde{x}_i^{\text{low}}(t)$ and $\tilde{x}_i^{\text{up}}(t)$.

By enlarging the set of the feasible charging profiles, for controllable customer $i \in \mathcal{N}_c$, the cumulative cost function value of the iterates that are generated by the iteration (7) is a lower bound for the cumulative cost function value of the iterates that are obtained by using the same iteration (7) without relaxation.

We now propose Algorithm 1 with the above mentioned set of relaxed feasible charging schedules $\tilde{\mathcal{F}}_i$, $i \in \mathcal{N}_c$ to compensate for the performance loss due to the presence of inelastic

customers. We consider that the controllable customers relax their charging constraints for the final J days.

Algorithm 1

Require: The distribution company knows \mathcal{N}_p , \mathcal{N}_l , \mathcal{N}_c , $\tilde{\mathcal{F}}_i$, $i \in \mathcal{N}_c$, J , and sets $\eta_u = 1/(2\sqrt{K})$. Each price-sensitive customer $i \in \mathcal{N}_p$ knows \mathcal{F}_i , and sets $\eta_i = 1/\sqrt{K}$. Each inelastic customer $i \in \mathcal{N}_l$ knows \mathcal{F}_i . Each controllable customer $i \in \mathcal{N}_c$ knows \mathcal{F}_i , $\tilde{\mathcal{F}}_i$, J , and sets $\eta_i = 1/\sqrt{K}$

- 1: Initialization: $k \leftarrow 1$, $x_i^k(t) \leftarrow S_i/T$, for all $i \in \mathcal{N}$, $t \in \mathcal{T}$
- 2: **for** $k = 1$ to K **do**
- 3: At the end of the k -th day, the distribution company gathers the charging profiles x_i^k , $i \in \mathcal{N}$, computes the prices $p_i^k(x^k) = (D^k + \sum_{i \in \mathcal{N}} x_i^k)$, $i \in \mathcal{N}$, notifies the prices to the price-sensitive customers and the controllable customers
- 4: At the end of the k -th day, every price-sensitive customer $i \in \mathcal{N}_p$ updates the charging profile by

$$h_i^{k+1} = h_i^k - \eta_i p_i^k, \\ x_i^{k+1} = \underset{x_i \in \mathcal{F}_i}{\operatorname{argmin}} \|x_i - h_i^{k+1}\|^2.$$

- 5: At the end of the k -th day, every inelastic customer $i \in \mathcal{N}_l$ updates the charging profile by

$$x_i^{k+1} = x_i^k.$$

- 6: At the end of the k -th day,
- 7: **if** $1 \leq k \leq (K - J)$ **then**
- 8: every controllable customer $i \in \mathcal{N}_c$ updates the charging profile by

$$h_i^{k+1} = h_i^k - \eta_i p_i^k, \\ x_i^{k+1} = \underset{x_i \in \mathcal{F}_i}{\operatorname{argmin}} \|x_i - h_i^{k+1}\|^2.$$

- 9: **else**

- 10: every controllable customer $i \in \mathcal{N}_c$ updates the charging profile by

$$h_i^{k+1} = h_i^k - \eta_i p_i^k, \\ \tilde{x}_i^{k+1} = \underset{\tilde{x}_i \in \tilde{\mathcal{F}}_i}{\operatorname{argmin}} \|\tilde{x}_i - h_i^{k+1}\|^2.$$

We have the following results to verify that Algorithm 1 is a no regret algorithm even with the presence of the inelastic customers.

Theorem IV.3. Consider that there are N EV customers out of which N_l are lazy and N_c are controllable. Assume that the controllable customers relax their charging constraints for the

final J days, $\eta_u = \frac{1}{2}\eta_i$, for all $i \in \mathcal{N}$, and the condition

$$\begin{aligned} & - \sum_{k=1}^{K-J} \left[\sum_{i \in \mathcal{N}_l} (x_i^k - x_i^*)^T (\epsilon_i^k) \right] \\ & + \sum_{k=K-J+1}^K \left[c_u^k(\tilde{x}^*) - c_u^k(x^*) - \sum_{i \in \mathcal{N}_l} (x_i^k - x_i^*)^T (\epsilon_i^k) \right] \leq 0, \end{aligned} \quad (28)$$

holds. Then the average regret of the distribution company is bounded as

$$\begin{aligned} \frac{R_u(K)}{K} & \leq \frac{1}{K} \left[\frac{1}{\eta_u} P_u + \frac{\eta_u}{2} \sum_{k=1}^{K-J} \left\| (\nabla c_u^k(x^k) + \epsilon^k) \right\|_*^2 \right] \\ & + \frac{1}{K} \left[\frac{1}{\eta_u} \tilde{P}_u + \frac{\eta_u}{2} \sum_{k=K-J+1}^K \left\| (\nabla c_u^k(\tilde{x}^k) + \epsilon^k) \right\|_*^2 \right], \end{aligned} \quad (29)$$

where

$$\tilde{P}_u := \max_{\tilde{x} \in \tilde{\mathcal{F}}} L(\tilde{x}) - \min_{\tilde{x} \in \tilde{\mathcal{F}}} L(\tilde{x}). \quad (30)$$

In particular, if $\eta = O(1/\sqrt{K})$, then the average regret converges to zero as $K \rightarrow \infty$.

The distribution company needs to design the value J and the set $\tilde{\mathcal{F}}_i$, $i \in \mathcal{N}_c$ to ensure that the condition (28) is satisfied. The condition (28) imposes requirements on the set of the relaxed feasible charging profiles and the number of days to relax the charging constraints for the controllable customers. If the base load remains the same from day to day and the distribution company knows the base load, the value $c_u^k(\tilde{x}^*)$ and $c_u^k(x^*)$, $k \in \mathbb{N}_{>0}$ can be computed, for example, using the algorithms proposed in [9], [10]. The distribution company is able to verify the condition (28) by using the upper bounds of $\|x_i^k\|$, $\|x_i^*\|$, and $\|\epsilon_i^k\|$, $i \in \mathcal{N}_l$, $k \in \mathbb{N}_{>0}$ as follows. Given the upper bounds of $\|x_i^k\|$, $\|x_i^*\|$, and $\|\epsilon_i^k\|$, $i \in \mathcal{N}_l$, $k \in \mathbb{N}_{>0}$, the distribution company requires that the value J and the set $\tilde{\mathcal{F}}_i$, $i \in \mathcal{N}_c$ satisfy the inequality

$$\begin{aligned} & \sum_{k=K-J+1}^K \left(c_u^k(x^*) - c_u^k(\tilde{x}^*) \right) \\ & \geq \sum_{k=K-J+1}^K \sum_{i \in \mathcal{N}_l} \left(\left(\|x_i^k\| + \|x_i^*\| \right) \|\epsilon_i^k\| \right) \\ & + \sum_{k=1}^{K-J} \sum_{i \in \mathcal{N}_l} \left(\left(\|x_i^k\| + \|x_i^*\| \right) \|\epsilon_i^k\| \right). \end{aligned} \quad (31)$$

V. NUMERICAL EXAMPLES

Assume that there are 20 customers. A time slot representing an interval of 30 minutes is used. There are $T = 24$ time slots. The starting time is set to 8:00 pm. For simplicity, we consider that all EV customers charge their EVs from the 9th to the 16th time slots¹. On the first day, the initial charging profiles are assumed to be uniformly distributed over the time slots. The maximum charging rate is set to $x_i^{\text{up}}(t) = 2$ kW, $i \in$

¹The problem formulation (3) allows EV customers to have different charging constraints. The initial charging time and final charging time mainly depend on the preferences of the EV customers.




Figure 1. Base load profile from 8:00 pm to 7:30 am (the next day).




Figure 2. Average regrets generated by OMD with and without prediction.

\mathcal{N} and the desired sum $S_i = 10$ kW, $i \in \mathcal{N}$. The simulation is carried out for total $K = 200$ days. We set the parameters $\eta_i = 0.05/\sqrt{K}$, $i \in \mathcal{N}$. We first examine the convergence of the static regret. The base load profile is given in Figure 1. Figure 2 shows the trajectories of the regrets with and without the prediction. The prediction M_i^k , $i \in \mathcal{N}$, $k \in \mathbb{N}_{>0}$ is set to $M_i^k(x_i^k) = \frac{1}{k-1} \sum_{\tilde{k}=1, \dots, k-1} \nabla c_i^{\tilde{k}}(x_i^{\tilde{k}})$, $i \in \mathcal{N}$. Figure 2 shows that the average regrets converge to zero and the average regret with the prediction converges faster than the one without the prediction. Figure 3 shows the static regrets with and without the prediction. Figure 3 shows that the regrets are sublinear functions of the number of days, which verifies the results in Proposition III.1.

We now consider a base load profile which does not remain the same from day to day. The base load is assumed to switch between two base load profiles (see Figure 4). We set the parameters $\eta_i = 0.005/\sqrt{K}$, $i \in \mathcal{N}$. Figure 5 shows the trajectories of the regrets with and without the prediction given the varying base load in Figure 4. Figure 5 shows that the average regrets converge to zero and the average regret with the prediction converges faster than the one without the prediction. Figure 6 shows the static regrets with and without the prediction. Figure 6 shows that the regrets are sublinear functions of the number of days. The results in Figure 5 and Figure 6 indicate that despite the fact that the base load is switching and the distribution company is not aware of this




Figure 3. Static regrets generated by OMD with and without prediction.




Figure 5. Average regrets generated by OMD with and without prediction given the varying base load profile in Figure 4.




Figure 4. Two different base load profiles. The actual base load is realized by switching between the two base load profiles from day to day.




Figure 6. Static regrets generated by OMD with and without prediction given the varying base load profile in Figure 4.

varying behavior of the base load, our algorithm still provides updates of the charging profiles having converging behavior of the regret. Average regret converges to zero means that in the long term, the average performance of the charging schedules x^k generated by the OMD algorithm approaches the performance obtained by x^* , where x^* solves (9).

Figure 7 provides the total load profiles at convergence with various number of inelastic customers. Figure 7 shows that as the total number of inelastic customers increases, the variation of the total load profile increases. The result indicates that the inelastic customers perturb the valley-filling load profile. Now suppose that there are 10 inelastic customers and 10 customers with relaxed constraints. We consider two different relaxation strategies. Relaxation 1 represents the strategy that allows the EVs to be charged over the entire time slots (rather than merely between the 9th and the 16th time slots), whereas Relaxation 2 represents the one extending the charging time slots to cover from the 8th to the 17th time slots. Figure 8 shows that as the allowed charging time slots are extended, the total load profile resembles a valley-filling profile, which verifies the effectiveness of the controllable customers to relief stress of the power grid.

VI. CONCLUSION

We have designed a framework for distributed charging control of EVs using online learning and online convex optimization. The proposed algorithm can be implemented without low-latency two-way communication between the distribution company and the EV customers, which fits in with the current communication infrastructure and protocols in the smart grid.

APPENDIX A PROOF OF PROPOSITION III.2

The update (7) yields

$$h_i^{k+1} = h_i^k - \eta_i \left(\sum_i x_i^k + D^k \right), \quad (32)$$

$$x_i^{k+1} = \underset{x_i \in \mathcal{F}_i}{\operatorname{argmin}} \eta_i x_i^T M_i^{k+1} + \|x_i - h_i^{k+1}\|^2, \quad (33)$$

where M_i^k is the prediction of the value $\sum_i x_i^k + D^k$.

The update (10) yields

$$\begin{bmatrix} h_1^{k+1} \\ h_2^{k+1} \\ \vdots \\ h_N^{k+1} \end{bmatrix}_{NT} = \begin{bmatrix} h_1^k \\ h_2^k \\ \vdots \\ h_N^k \end{bmatrix}_{NT} - \eta_u \begin{bmatrix} 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \\ 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \\ \vdots \\ 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \end{bmatrix}_{NT}, \quad (34)$$




Figure 7. The total load profiles with various numbers of inelastic customers, the optimal total load profile, and the base load profile.




Figure 8. The total load profile generated by OMD, the optimal total load profile, base load profile, and the total load profile with the relaxations.

$$x^{k+1} = \underset{x \in \mathcal{F}}{\operatorname{argmin}} \eta_u x^T M_u^{k+1} + \|x - h_u^{k+1}\|^2, \quad (35)$$

where $h_u^{k+1} = [h_1^{k+1}, \dots, h_N^{k+1}]^T$ and M_u^k is the prediction of

$$\begin{bmatrix} 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \\ 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \\ \vdots \\ 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \end{bmatrix}_{NT}. \quad (36)$$

By comparing (34), (35) with (32), (33) and substituting $\eta_u = \frac{1}{2}\eta_i$, $i \in \mathcal{N}$, we have that the updates (35) and (33) are identical. Since the updates coincide, the average regret of the distribution company converges to zero as $k \rightarrow \infty$.

APPENDIX B PROOF OF THEOREM IV.1

The proof technique that we use to derive the tracking regret bound in (19) is similar to the one used in [12, Theorem 4]. The main step is to bound the difference $c_u^k(x^k) - c_u^k(x^{k*})$ instead of the difference $c_u^k(x^k) - c_u^k(x^*)$ that is considered in the static regret (13). However, in [12], the authors derive the tracking regret bounds for a different regret minimization algorithm rather than OMD.

Following the proof of [23, Lemma 2], we have

$$c_u^k(x^k) - c_u^k(x^{k*}) \leq (x^k - x^{k*})^T \nabla c_u^k(x^k), \quad (37)$$

and

$$\begin{aligned} (x^k - x^{k*})^T \nabla c_u^k(x^k) &\leq \frac{\eta_u}{2} \|\nabla c_u^k(x^k) - M_u^k\|_*^2 \\ &+ \frac{1}{\eta_u} (D_{L_u}(x^{k*}, h_u^k) - D_{L_u}(x^{k*}, h_u^{k+1})). \end{aligned} \quad (38)$$

Furthermore,

$$\begin{aligned} D_{L_u}(x^{k*}, h_u^k) - D_{L_u}(x^{k*}, h_u^{k+1}) &= L_u(x^{k*}) - L_u(h_u^k) - \nabla L_u(h_u^k)^T (x^{k*} - h_u^k) \\ &- L_u(x^{k*}) + L_u(h_u^{k+1}) + \nabla L_u(h_u^{k+1})^T (x^{k*} - h_u^{k+1}) \\ &= L_u(h_u^{k+1}) - L_u(h_u^k) + \nabla L_u(h_u^{k+1})^T (x^{k+1*} - h_u^{k+1}) \\ &- \nabla L_u(h_u^k)^T (x^{k*} - h_u^k) - \nabla L_u(h_u^{k+1})^T (x^{k+1*} - x^{k*}). \end{aligned} \quad (39)$$

The remainder of the proof is followed by summing over $k = 1, \dots, K$ and collecting terms.

APPENDIX C PROOF OF THEOREM IV.2

Following the proof of [23, Lemma 2], we have

$$c_u^k(x^k) - c_u^k(x^*) \leq (x^k - x^*)^T \nabla c_u^k(x^k), \quad (40)$$

and

$$\begin{aligned} (x^k - x^*)^T \nabla c_u^k(x^k) &\leq \frac{\eta_u}{2} \|\nabla c_u^k(x^k)\|_*^2 \\ &+ \frac{1}{\eta_u} (D_{L_u}(x^*, h_u^k) - D_{L_u}(x^*, h_u^{k+1})), \end{aligned} \quad (41)$$

where

$$\nabla c_u^k(x^k) = \begin{bmatrix} 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \\ \vdots \\ 2(D^k + \sum_{i \in \mathcal{N}} x_i^k) \end{bmatrix}_{NT}. \quad (42)$$

For each $i \in \mathcal{N}_l$, the cost function c_i^k is selected as a constant function noted in Section IV-B. The corresponding gradient of the lazy customer's cost function is zero, namely, for lazy customer $i \in \mathcal{N}_l$,

$$\nabla c_i^k(x^k) = \left(D^k + \sum_{i \in \mathcal{N}} x_i^k \right) + \epsilon_i^k = 0, \quad (43)$$

where $\epsilon_i^k = -(D^k + \sum_{i \in \mathcal{N}} x_i^k)$. Move ϵ_i^k , $i \in \mathcal{N}_l$ to the right hand side of the inequality in (41). The remainder of the proof is followed by summing over $k = 1, \dots, K$ and collecting terms.

APPENDIX D PROOF OF THEOREM IV.3

First observe that, for $i \in \mathcal{N}_l$,

$$c_u^k(\tilde{x}^k) - c_u^k(x^*) = [c_u^k(\tilde{x}^k) - c_u^k(\tilde{x}^*)] + [c_u^k(\tilde{x}^*) - c_u^k(x^*)]. \quad (44)$$

Following the proof of [23, Lemma 2],

$$\begin{aligned}
c_u^k(\tilde{x}^k) - c_u^k(\tilde{x}^*) &\leq (\tilde{x}^k - \tilde{x}^*)^T \nabla c_u^k(\tilde{x}^k) \\
&\leq \frac{\eta_u}{2} \|\nabla c_u^k(\tilde{x}^k)\|_*^2 \\
&\quad + \frac{1}{\eta_u} (D_{L_u}(\tilde{x}^*, h_u^k) - D_{L_u}(\tilde{x}^*, h_u^{k+1})).
\end{aligned} \tag{45}$$

The remainder of the proof is followed by summing over $k = 1, \dots, K$ and collecting terms.

REFERENCES

- [1] Con Edison, “Demand response/distribution load relief program,” 2012, <http://www.coned.com/energyefficiency>.
- [2] Office of the General Counsel, U.S. Department of Energy, “Communications requirements of smart grid technology,” <http://energy.gov/gc/office-general-counsel>.
- [3] Wisconsin Public Service, “Contracted direct load control,” 2012, <http://www.wisconsinpublicservice.com/business>.
- [4] A. Blum and Y. Mansour, “From external to internal regret,” *J. of Machine Learning Research*, vol. 8, pp. 1307-1324, 2007.
- [5] M. C. Caramanis and J. M. Foster, “Coupling of day ahead and real-time power markets for energy and reserves incorporating local distribution network costs and congestion,” in *Proc. of 2010 IEEE Allerton Conference*, pp. 42-49, 2010.
- [6] N. Cesa-Bianchi and G. Lugosi, “Prediction, Learning, and Games,” Cambridge University Press, 2006.
- [7] L. Chen, N. Li, L. Jiang, and S. H. Low, “Optimal demand response: problem formulation and deterministic case,” *Control and Optimization Methods for Electric Smart Grids, Power Electronics and Power Systems*, vol. 3, pp. 63-85, 2012.
- [8] K. Clement, E. Haesen, and J. Driesen, “Coordinated charging of multiple plug-in hybrid electric vehicles in residential distribution grids,” in *Proc. of IEEE/PES Power Systems Conference and Exposition*, pp. 1-7, 2009.
- [9] L. Gan, U. Topcu, and S. Low, “Optimal decentralized protocol for electric vehicle charging,” *IEEE Trans. on Power Systems*, vol. 28, no. 2, pp. 940-951, 2013.
- [10] L. Gan, A. Wierman, U. Topcu, U. Chen, and S. H. Low, “Real-time deferrable load control: handling the uncertainties of renewable generation,” in *Proc. of International Conference on Future Energy Systems*, pp. 113-124, 2013.
- [11] G. B. Giannakis, V. Kekatos, N. Gatsis, S.-J. Kim, H. Zhu, and B. Wollenberg, “Monitoring and optimization for power grids: a signal processing perspective,” *IEEE Signal Processing Magazine*, vol. 30, no. 5, pp. 107-128, Sep. 2013.
- [12] E. C. Hall and R. M. Willett, “Dynamical Models and tracking regret in online convex programming,” in *Proc. of the International Conference on Machine Learning*, pp. 579-587, 2013.
- [13] E. Hazan, A. Agarwal, and S. Kale, “Logarithmic regret algorithms for online convex optimization,” *Machine Learning Journal*, vol. 69(2-3), pp. 169-192, 2007.
- [14] L. Jia, L. Tong, and Q. Zhao, “Retail pricing for stochastic demand with unknown parameters: an online machine learning approach,” in *Proc. of 2013 IEEE Allerton Conference*, pp. 1353-1358, 2013.
- [15] L. Kelly, A. Rowe, and P. Wild, “Analyzing the impacts of plug-in electric vehicles on distribution networks in British Columbia,” in *Proc. of IEEE Electrical Power and Energy Conference*, pp. 1-6, 2009.
- [16] S.-J. Kim and G. B. Giannakis, “Real-time electricity pricing for demand response using online convex optimization,” in *Proc. of IEEE Innovative Smart Grid Technologies Conference*, pp. 19-22, 2014.
- [17] P. Kumar, J. S. B. Mitchell, and E. A. Yildirim, “Approximate minimum enclosing balls in high dimensions using core-sets,” *J. of Experimental Algorithmics*, vol. 8, no. 1.1, pp. 1-29, 2003.
- [18] J. A. P. Lopes, F. J. Soares, and P. M. R. Almeida, “Integration of electric vehicles in the electric power system,” in *Proc. of the IEEE*, vol. 99, no. 1, pp. 168-183, 2011.
- [19] W.-J. Ma, V. Gupta, and U. Topcu, “On distributed charging control of electric vehicle with power network capacity constraints,” in *Proc. of IEEE American Control Conference*, pp. 4306-4311, 2014.
- [20] Z. Ma, D. S. Callaway, and I. A. Hiskens, “Optimal charging control for plug-in electric vehicles,” *Control and Optimization Methods for Electric Smart Grids*, Springer, 2012.
- [21] Z. Ma, D. S. Callaway, and I. A. Hiskens, “Decentralized charging control of large populations of plug-in electric vehicles,” in *Proc. of IEEE Conference on Decision and Control*, pp. 206-212, 2010.
- [22] F. Rahimi and A. Ipakchi, “Demand response as a market resource under the smart grid paradigm,” *IEEE Trans. on Smart Grid*, vol. 1, no. 1, pp. 82-88, 2010.
- [23] A. Rakhlin and K. Sridharan, “Online learning with predictable sequences,” *J. of Machine Learning Research: Workshop and Conference Proceedings*, vol. 30, pp. 1-27, 2013.
- [24] S. Shalev-Shwartz, “Online Learning and Online Convex Optimization,” *Foundations and Trends in Machine Learning*, vol. 4, no. 2, pp. 107-194, 2012.
- [25] N. Y. Soltani, S.-J. Kim, and G. B. Giannakis, “Online learning of load elasticity for electric vehicle charging,” in *Proc. of IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing*, pp. 15-18, 2013.
- [26] E. Sortomme, M. M. Hindi, S. D. J. MacPherson, and S. S. Venkata, “Coordinated charging of plug-in hybrid electric vehicles to minimize distribution system losses,” *IEEE Trans. on Smart Grid*, vol. 2, no. 1, pp. 198-205, 2011.
- [27] Martin Zinkevich, “Online convex programming and generalized infinitesimal gradient ascent,” in *Proc. of International Conference on Machine Learning*, pp. 928-936, 2003.